

The web after websites: our vision for the agentic web

By Léon Missoul, August 2026

Sometime this spring, machines quietly became the majority of the web. Most sites have not noticed, because they still look the way they always did: pages, written for human eyes, waiting for a visit that increasingly never comes. We think the visit is going away entirely, and faster than almost anyone building websites is ready for.

Five predictions

A vague vision costs nothing, so we are putting dates on ours.

1. **By mid-2027, the majority of first contact with a B2B service firm happens without a human ever loading its homepage.** By the time anyone reaches out, an assistant has already read the site, answered the question, and moved on.
2. **By end of 2027, an agent can find a site's action endpoint and connect to it with no one wiring it up by hand.** This is the hinge. Nothing else on this list happens until it does.
3. **By mid-2028, organic search sends B2B service firms less than a third of the traffic it sent in 2024.** On the research questions buyers actually type, comparisons and how-tos, the click is all but gone: the assistant answers, names a few sources, and sends no one anywhere.
4. **By 2029, a real share of quotes, bookings, and orders in professional services begin agent-to-site, with no page view anywhere in the transaction.** The first thing a firm learns about the deal is that it already started.
5. **By end of 2029, new commercial sites take more agent traffic through structured endpoints than through browsers imitating a human.** The bridge starts giving way to the road it was built for. As far as we can tell, nobody else has put a date on that.

These come from our own research and from what we watch happen across the B2B sites we build and operate. Some of them will be wrong. That is what a date is for. **Hold us to them.**

Three things that already broke

Machines became the majority of the web this year. In June 2026, Cloudflare reported that automated systems generate 57.5% of the HTTP requests on its network, against 42.5% from humans, the first such crossover ever recorded. Its own CEO had called the moment coming and still missed it by 18 months, overtaken by agentic traffic that grew roughly 7,851% year over year. The asymmetry explains the speed. A human shopping visits perhaps five sites. An agent running the same errand, by Prince's own illustration, visits five thousand. Multiply that across every errand on earth and the crossover stops being surprising; the surprise is that it took until 2026.

The second one is closer to home. Rankings still work. The clicks they used to deliver do not. Google now answers a growing share of queries on the results page itself, assistants answer them before a results page even loads, and a first-place ranking increasingly buys a citation instead of a visit. Publishers watching their server logs see it at industrial scale: crawlers pull thousands of pages for every referral they hand back. The old bargain that paid for the open web, content traded for visits, is being rewritten by one side of the table.

The third one is ours. For a handful of the sites we run, we track whether assistants name the business when a buyer asks the question it sells into. One now surfaces in roughly one ChatGPT answer in three. The visibility is real, and it is winnable, and it sends almost nobody. Four months of being named that way, across the whole portfolio, came to 322 visitors, under half a percent of human traffic, and the number sat still the whole time: 0.63% of visits in May, 0.28% in June, 0.35% in July, 0.29% in August. The reports all call this traffic explosive. On our servers it never moved.

The few who come are worth watching. An assistant sends us the most engaged visitors of any channel, roughly half of them engaging with the page against 38% across the portfolio, and the gap holds in every window we measure. The assistant did the qualifying already. It answered the question itself and forwarded only the people who still had a reason to show up.

Put those together and the pattern is simple: the reader of the web changed before the web itself actually did. That is what a transition looks like from the inside.

Search changed masters

Ask an AI assistant a question and it searches. Behind most answers sits an ordinary query to Google or Bing; the assistant reads what ranks and answers from it. So ranking still matters, but it now buys something it never sold before: being the source the machine reads first, cites first, trusts first. The work is the same. The reader is a machine.

That reader has different taste. Your hero animation means nothing to it. Your art direction means nothing to it. What it wants is the answer to a specific question, stated plainly, in a structure it can parse without guessing. Most B2B sites are autobiographies. Almost none answer questions, which makes answering them the cheapest competitive advantage on the internet right now.

The arrangement is also temporary, and that is the part most people miss. Search engines exist to help humans find pages. Agents lean on them only until something better arrives,

and when it does, the detour through search collapses. Every euro you spend ranking today is a loan against the next web. It pays out only if what you publish is worth reading to a machine.

The catalog replaces the index

There are two ways this could go, and most of the industry is building for one of them.

The first: agents keep using the human web, opening a browser, rendering pages, clicking buttons the way a person would. Several major labs ship exactly this today, and it works. Enterprises run it in production at real scale, mostly against decades-old portals and legacy software that will never expose a modern interface. Hold onto that detail.

We are building for the second: the web grows a machine layer. Sites expose structured endpoints, agents call them directly, and nothing renders, because nothing needs to.

Our reason for betting against the browser agent has nothing to do with whether it works. It works the way a horseless carriage worked. The first cars were shaped like carriages because the carriage was the only form anyone knew; the first mobile sites were desktop pages shrunk to fit. A browser agent imitates a human for the same reason, because the human interface is the only door most sites leave open, so it renders pixels, reads layouts, and clicks through forms to reach information a machine could have asked for in a single call. Inside enterprises it thrives exactly where no better path exists and none is coming; a government portal from 2003 is never going to ship an endpoint. The public web runs the other way. There, every browser agent is proof of demand and proof of waste at once, and the day a site offers an endpoint, agents stop rendering its pages, for the same reason developers stop scraping a site the moment it offers an API. The browser agent is the bridge that makes its own destination inevitable.

That destination is close, because the hard part already shipped. The Model Context Protocol gives AI systems a standard way to use external tools; it runs in production across the major AI products, and thousands of companies operate MCP servers today. The interaction layer is built. The missing piece is discovery. Connecting an endpoint to an AI tool is still a manual act, and an agent has no way to stumble onto your MCP server the way a browser stumbles onto your homepage.

The web has solved this exact problem once before. DNS, links, and crawlers turned pages into something anyone could find; the same pattern pointed at capabilities turns endpoints into something any agent can find. The first attempts are already live. In June 2026, Google, Microsoft, and nine other companies published Agentic Resource Discovery, or ARD: a site lists its machine interfaces in a single catalog file at a fixed path on its own domain, and registries crawl those catalogs the way search engines crawl pages. A W3C community group led by Google and Microsoft engineers is drafting WebMCP alongside it, which lets a site register callable tools directly in the page. Both are weeks old, and adoption is near zero; a June census found that not even the companies behind ARD were serving a catalog file yet. Read that as the ground floor of a standard, not as a rebuttal, and as the reason the window to be early is still open. Discovery is the contested layer, and whoever wins it owns the next web.

Once discovery works at scale, the shape of a website turns inside out. Today a site is a stack of pages a human reads and a machine scrapes. Tomorrow it is a set of capabilities a machine calls: answer this question, check this availability, produce this quote, book this meeting, place this order, put these two people in touch. The human pages stay, the way phone lines stayed, but they stop being the main way in. Search engines rank pages. The next layer catalogs capabilities. That is the whole shift in one line.

The early failures make the case better than the successes do. OpenAI launched Instant Checkout in September 2025 and pulled it roughly five months later with almost nothing to show: tax handling unsolved, inventory out of sync, fewer than fifteen merchants live out of millions eligible, and shoppers who researched inside the assistant but paid in the checkout they already trusted. The interface failed. The layer underneath it held. Shopify switched agent-facing storefronts on by default for every eligible store in March 2026, after AI-attributed orders rose elevenfold through late 2025; by the first quarter of 2026 its AI-driven traffic was up eightfold year over year and orders from AI-powered search nearly thirteenfold. The payment protocols consolidated under industry bodies instead of vanishing, with Google handing its agent payments protocol to the FIDO Alliance this spring alongside some sixty other organizations. The early web went the same way. Push channels died. HTTP won.

The demand is already here; only the rails are late. Ninety-four percent of B2B buyers now use generative AI as a research tool, and most say they would place large orders through self-service with no sales call. What happens next is genuinely contested. Gartner projects fifteen trillion dollars in B2B purchases flowing through agents by 2028; soberer forecasts land two orders of magnitude lower by the end of the decade. The direction is unanimous and the speed is anyone's guess, which is why we put a date on the behavior and not on the size. As for the browser agent, it is not fading yet, it is about to get more common: no mainstream agent calls a site's structured tools today, and Google is expected to ship an agentic browser mode by default before the year is out, which multiplies that traffic overnight. Our bet is not that browser agents disappear. It is that they are the top of an approach rather than the start of one, and that the structured path passes them on new commercial sites once discovery is solved.

One caveat, usually skipped. The destination is clear to us; the road is not. Discovery still runs through search today, and search belongs to two companies. How fast the catalog layer arrives comes down to how Google, and to a lesser extent Microsoft and the AI labs, decide to play their hands. Anyone promising you a certain timeline is selling something. We will give you the direction and take the risk on the dates.

What this means for a business

No checklist. No file you upload to fix it. Anyone selling a one-step cure for the agentic web is selling snake oil. Here is what we think holds no matter which protocol wins, and none of it is exotic.

Know exactly who you serve and the question they arrive with. An agent carries a real buyer's real constraints, and a site pitched at a vague audience has nothing to say to a specific question. Context is the asset. The rest is rendering.

Build the site to answer questions, not to describe you. The next reader is running an errand for someone. It needs your pricing logic, your service boundaries, your availability, your proof, laid out plainly and found without interpretation.

One of our clients runs Meta ads for a single kind of business: high-ticket home improvement, the contractors and custom-build firms where the average job is large and trust decides everything. Because they know their buyer that precisely, we could build the site around the questions that buyer actually asks, backed with the proof those buyers look for, instead of another page describing the agency. There was nothing clever about it. Now when someone asks an assistant for the best Meta ads agency for contractors in their area, that firm comes up. A clear foundation and the right context were the whole trick.

Keep the content legible to a machine. Clear structure, real answers in real sentences, nothing important trapped inside an image or an animation. When a capable reader cannot lift your offer off the page in one pass, the problem was never the content. It was the structure.

Treat your web presence as infrastructure, not a project. The four-year rebuild assumes a web that sits still, and this one does not. A site now has to be kept current against a moving target: new formats, new readers, new interfaces as the catalog layer settles. The firms that win a transition are rarely the ones who called it. They are the ones standing in the right place when it arrives.

Read those four again. None of it is futurism. It is what a good website should always have been. The agentic web does not change what quality is. It just takes away the option of faking it.

Our position

We build websites for a living, and here we are arguing that the website as everyone knows it is a passing format. We know how that sounds. A firm that lives on four-year rebuilds has every reason to keep quiet about this, and instead we are saying it out loud, because we read the traffic across our clients every week and the numbers stopped telling the old story.

So we build for the reader we think is coming. Structured context first, rendered pages second, and Orbit underneath to keep that structure current as the standards move, because a site built for this shift is not a thing you ship once. It is a thing you maintain toward a target that keeps moving.

We published five dated predictions at the top of this paper. Twelve months from now we will come back and score them in the open. If we called it wrong, you will see exactly where.

Transitions never feel like transitions from the inside. The homepage still loads, the rankings still update, the analytics still fill with sessions, and everything looks like it always has. Meanwhile the majority of the traffic is already machines, reading a web that was never built for them, waiting for one missing piece. The last shift of this size took the storefront and put it

on a screen. This one takes the screen out of the transaction entirely. **We would rather be early and wrong on the dates than late and right on the direction.**

Notes

- Cloudflare / Matthew Prince, June 3, 2026: bots pass human traffic for the first time, 57.5% vs 42.5% of HTTP requests to HTML content; crossover ~18 months ahead of his own SXSW prediction. Tom's Hardware: <https://www.tomshardware.com/tech-industry/artificial-intelligence/bots-have-now-passed-human-traffic-online-cloudflare-boss-laments-says-agentic-traffic-wasnt-expected-to-eclipse-real-people-until-next-year> and Search Engine Land: <https://searchengineland.com/cloudflare-bots-webpage-requests-479608>
- Agentic traffic growth ~7,851% year over year; five-vs-5,000 asymmetry per Prince at SXSW: <https://ppc.land/us-sends-53-5-of-global-bot-traffic-decodo-analysis-finds/> and https://ground.news/article/bots-now-generate-more-web-traffic-than-humans-says-cloudflare_270264
- Agentic Resource Discovery (ARD) specification, ai-catalog.json at a well-known path, published June 17, 2026 by Google, Microsoft, GitHub, Hugging Face, Cisco, Databricks, GoDaddy, NVIDIA, Salesforce, ServiceNow, and Snowflake: <https://developers.googleblog.com/announcing-the-agentic-resource-discovery-specification/> and <https://www.infoq.com/news/2026/07/agentic-resource-discovery-spec/> ; June 2026 adoption census finding zero discoverable catalogs among 39 major sites including the working-group members: <https://www.synscribe.com/agentic-discovery/agentic-resource-discovery>
- WebMCP, W3C Web Machine Learning Community Group draft (Google and Microsoft editors): <https://webmachinelearning.github.io/webmcp/>
- OpenAI Instant Checkout: launched September 29, 2025, pulled March 4, 2026, fewer than 15 merchants live, unsolved tax and inventory handling. Forbes: <https://www.forbes.com/sites/jasongoldberg/2026/03/10/why-openais-checkout-retreat-spells-trouble-for-its-commerce-strategy/>
- Shopify Agentic Storefronts on by default for eligible US merchants, March 2026: <https://aiadvantageagency.com/shopify-agentic-storefronts/>
- Shopify AI figures: AI traffic up 7x and AI-attributed orders up 11x, January to November 2025 (Q3 2025 earnings), TechCrunch: <https://techcrunch.com/2025/11/04/shopify-says-ai-traffic-is-up-7x-since-january-ai-driven-orders-are-up-11x/> ; Q1 2026: AI traffic up 8x, AI-search orders up nearly 13x year over year: <https://finance.yahoo.com/markets/stocks/articles/shopifys-ai-push-sustain-more-161700877.html>
- Google donates Agent Payments Protocol (AP2) to the FIDO Alliance, April 28, 2026, with sixty contributing organizations: <https://blog.google/products-and-platforms/platforms/google-pay/agent-payments-protocol-fido-alliance/> and <https://nohacks.co/blog/agent-payments-protocol-60-organizations>

- Additional context: The Economist, interview with Elon Musk, recorded July 20, 2026 at Giga Texas, released July 23, 2026:
<https://www.youtube.com/watch?v=8TjK-s0468w>
- Luniq first-party measurement, 15 August 2026. Across fifteen small Belgian business sites we operate and instrument with our own analytics, 29 April to 15 August 2026: AI assistants referred 322 human visitors, 0.38% of 83,526 human sessions, across ten of the fifteen sites (ChatGPT 239, Claude 29, Perplexity 25, Gemini 22, Copilot 6, Poe 1). The monthly share did not rise over the period (0.63% in May, 0.28% in June, 0.35% in July, 0.29% in August), and detection was widened on 5 August 2026, so the flat trend is not an artifact of weaker earlier measurement. Assistant-referred sessions engaged at 50.3% against a 38.5% portfolio baseline (two-proportion $z = 4.35$, $p < 0.0001$). Answer-surface tracking on one site recorded citation in 22 of 75 ChatGPT responses to buyer-intent prompts between 28 July and 11 August 2026. Two limits worth stating plainly: the measurement is a browser-executed script that sees only visits which render the page and run JavaScript, so it records nothing about crawlers or server-log traffic and no figure here describes crawler activity; and the instrument itself is young, live roughly four months, so these numbers may still carry gaps or errors we have not yet caught. We separately tested whether click-through at constant ranking position has fallen year over year across this portfolio and found no effect that survives controlling for query composition and seasonality; with only five sites carrying a full year of Search Console history, the sample is too small to answer that question, and we will re-run it in Q1 2027.